Noname manuscript No.
(will be inserted by the editor)

A penalized approach to the bivariate logistic regression model for the association between ordinal responses

Marco Enea · **Gianfranco Lovison**

arXiv:1407.1751v1 [math.ST] 7 Jul 2014

Received: date / Accepted: date

Abstract Bivariate ordered logistic models (BOLMs) are appealing to jointly model the marginal distribution of two ordered responses and their association, given a set of covariates. When the number of categories of the responses increases, the number of global odds ratios (or their reparametrizations) to be estimated also increases and estimating the association structure becomes crucial for this type of data. In fact, such data could be too "rich" to be fully modelled with an ordinary BOLM while, sometimes, the well-known Dale's model could be too parsimonious to provide a good fit. In addition, when the cross-tabulation of the responses contains some zeros, for a number of model configurations, including the bivariate version of the partial proportional odds model (PPOM), estimation of a BOLM by the Fisher-scoring algorithm may either fail or estimate a too "irregular" association structure.

In this work, we propose to use a nonparametric approach for the maximum likelihood estimation of a BOLM. We apply penalties to the differences between adjacent row and column effects. As a result, estimation is less demanding than an ordinary BOLM, permitting the fit of PPOMs and/or the smoothing of the marginal and association parameters by polynomial curves and surfaces, with scores chosen by the data. Model selection is based on the penalized log-likelihood ratio, whose limiting distribution has been studied through simulations, and AIC. Our proposal is compared to the Goodman's model and the Dale's model, in terms of goodness-of-fit and parsimony, on a literature data set. Finally, an application on an original data set of liver disease patients is proposed.

Marco Enea · Gianfranco Lovison
Dipartimento di Scienze Economiche Aziendali e Statistiche, University of Palermo, Palermo, Italy;
Tel.: +39-09123895334
Fax: +39-091485726
E-mail: marco.enea@unipa.it; gianfranco.lovison@unipa.it

Keywords Dale model · bivariate ordered logistic model · penalized maximum likelihood estimation · ordinal association

1 Introduction

Models for association play a central role in ordered categorical data analysis. For the multivariate case, marginal models (MMs) represent a natural choice to model marginal distributions of the responses given covariates. An example of full likelihood based marginal model is Dale (1986). A similar model, the multivariate logistic model described in Glonek and McCullagh (1995), but restricted to the bivariate ordered version, is the basis on which we develop our proposal. Some open, or at least not completely solved, problems about estimation of a multivariate ordered logistic model are of computational type and concern maximum likelihood (ML) estimation by iterative algorithms, often providing invalid estimates at the kth step, exceeding the boundaries of the parameter space. Some of such problems could be solved as in Colombi and Forcina (2001) and Bartolucci and Forcina (2002) by including strict inequality constraints. However, constrained ML estimation is appealing only when a particular application implies natural ordering constraints. On the contrary, when the ordering is not fully reliable, or externally imposed, like in responses which arise from discretized versions of latent continuous variables, using inequality constraints may not be appropriate. Indeed, due to lack of subject-matter knowledge that yields natural restrictions on marginal distributions, no strict ordering constraints are appropriate, and more helpful and flexible approaches are necessary. In these situations, a nonparametric approach may be useful (Dardanoni and Forcina, 1998). Within the possible range of nonparametric approaches, penalization is the one con-

sidered in this paper. Surprisingly, there is little literature on penalization applied to marginal models. Desantis et al (2008) apply a ridge penalty to a latent class model for ordinal data to stabilize ML estimation, that would otherwise not be computationally feasible without application of strict constraints. Other contributions deal mainly with forms of longitudinal (Gieger, 1997; Fahrmeier et al, 1999) or horizontal (Bustami et al, 2001) nonparametric modeling. The former focuses on smoothing of variation of marginal and association parameters over time, the latter refers to a form of smoothing on covariates, often by using splines.

Our proposal is based on a form of vertical smoothing - that is across response levels - of the regression parameters in order to regularize the parameter space and/or fit polynomial models using scores "chosen by the data". After recalling the Dale, the Gloneck-McCullagh and the bivariate partial proportional odds models in Section 2, and the penalized ML estimation approach in Section 3, the penalty terms we propose to use are introduced in Section 3. By simulation, we show the advantage of using the proposed approach in Section 4 and check the asymptotic behaviour of the penalized deviance statistic in Section 5. Two applications are considered in Section 6: in the first one, we compare our proposal to the Dale (1986) and the Goodman (1979) models on a literature data set, whereas the second one is about a data set of liver disease patients.

2 Bivariate ordered logit models

For two ordered outcomes A_1 and A_2, define the row and column marginal cumulative probabilities of a $D_1 \times D_2$ contingency table A_1A_2 as

$$\mu_{r\cdot} = P(A_1 \le r) = \sum_{i \le r} \pi_{i\cdot}, \quad \mu_{\cdot c} = P(A_2 \le c) = \sum_{j \le c} \pi_{\cdot j},$$

and the upper-left quadrant probabilities as

$$\mu_{rc} = P(A_1 \le r, A_2 \le c) = \sum_{i \le r} \sum_{j \le c} \pi_{ij},$$

with $r = 1, \ldots, D_1$, $c = 1, \ldots, D_2$. By differencing we obtain

$$P(A_1 \le r, A_2 > c) = \mu_{r\cdot} - \mu_{rc},$$
$$P(A_1 > r, A_2 \le c) = \mu_{\cdot c} - \mu_{rc},$$
$$P(A_1 > r, A_2 > c) = 1 - \mu_{r\cdot} - \mu_{\cdot c} + \mu_{rc}.$$

By choosing the cumulative odds as ordinal risk measures, and the logit as link function, we obtain the *global logits* (or *log global odds*):

$$\log \phi_{1r} = \text{logit}[P(A_1 \le r)] = \log(\mu_{r\cdot}) - \log(1 - \mu_{r\cdot}), \qquad (1)$$

$$\log \phi_{2c} = \text{logit}[P(A_2 \le c)] = \log(\mu_{\cdot c}) - \log(1 - \mu_{\cdot c}), \qquad (2)$$

$r = 1, \ldots, D_1 - 1$, $c = 1, \ldots, D_2 - 1$. By choosing the cross-products of quadrant probabilities as ordinal association measures, and the natural logarithm as link function, the *log global odds ratios* (or log-GORs) are defined as:

$$\begin{aligned} \log \psi_{rc} &= \log \frac{P(A_1 \le r, A_2 \le c) P(A_1 > r, A_2 > c)}{P(A_1 \le r, A_2 > c) P(A_1 > r, A_2 \le c)} \\ &= \log \frac{\mu_{rc}(1 - \mu_{r\cdot} - \mu_{\cdot c} + \mu_{rc})}{(\mu_{r\cdot} - \mu_{rc})(\mu_{\cdot c} - \mu_{rc})}. \end{aligned} \qquad (3)$$

Given the three parameters $\mu_{r\cdot}$, $\mu_{\cdot c}$, and ψ_{rc}, we may find the corresponding joint cumulative probabilities with the following inversion formula:

$$\mu_{rc} = \begin{cases} \frac{1}{2}(\psi_{rc} - 1)^{-1}(a_{rc} - \sqrt{a_{rc}^2 + b_{rc}}) & \text{if } \psi \ne 1, \\ \mu_{r\cdot}\mu_{\cdot c} & \text{if } \psi = 1, \end{cases} \qquad (4)$$

where $a = 1 + (\mu_{r\cdot} + \mu_{\cdot c})(\psi_{rc} - 1)$ and $b = -4\psi_{rc}(\psi_{rc} - 1)\mu_{r\cdot}\mu_{\cdot c}$. If the cumulative probabilities $\mu_{r\cdot}$ and $\mu_{\cdot c}$ satisfy the constraints $\mu_{r\cdot} < \mu_{r+1,\cdot}$ for $r = 1, \ldots, D_1 - 1$, and $\mu_{\cdot c} < \mu_{\cdot,c+1}$ for $c = 1, \ldots, D_2 - 1$, and the global odds ratios are not dependent on the category, that is $\psi_{rc} = \psi$, then (4) is a Plackett distribution (Plackett, 1965). Thus, the bivariate Dale regression model for $(\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \boldsymbol{\psi}_{12})'$ is as follows:

$$\begin{cases} \log[\phi_{1,r}(\boldsymbol{x})] = \beta_{10r} - \boldsymbol{\beta}_1' \boldsymbol{x}, \\ \log[\phi_{2,c}(\boldsymbol{x})] = \beta_{20c} - \boldsymbol{\beta}_2' \boldsymbol{x}, \\ \log[\psi_{rc}(\boldsymbol{x})] = \alpha + \rho_{1r} + \rho_{2c} + \sigma_{rc} - \boldsymbol{\beta}_3' \boldsymbol{x}, \end{cases} \qquad (5)$$

$r = 1, ..., D_1 - 1$, $c = 1, ..., D_2 - 1$. This model does not require marginal scores for responses and it is also invariant under any monotonic transformation of the marginal responses. Further, since the model is based on global odds ratios, collapsing adjacent row or column categories does not produce any effect in parameter interpretation, which remains unchanged with the exception of the intercepts related to the collapsed categories. This is in contrast with the RC Goodman model which uses local cross-ratios. In a more general framework than (5), Glonek and McCullagh (1995) introduce the *multivariate logistic model*:

$$\mathbf{C}' \log(\mathbf{L}\boldsymbol{\pi}) = \mathbf{X}\boldsymbol{\beta}, \qquad (6)$$

where $\mathbf{C}$ is a contrasts matrix, $\mathbf{L}$ is a matrix with elements $a_{ij} \ge 0$ such that $\mathbf{L}\boldsymbol{\pi} = \boldsymbol{\mu}$, $\boldsymbol{\eta} = \mathbf{C}' \log(\mathbf{L}\boldsymbol{\pi})$ is the parameter vector of interest, and $\mathbf{X}$, an $n \times p$ matrix, with $n = \prod_{k=1}^{K} D_k$. Although formulation (6) is referred to $K \ge 2$ responses, here only two responses A_1 and A_2 are considered. The components of $\mathbf{C}' \log(\mathbf{L}\boldsymbol{\pi})$ are symbolically denoted by $\boldsymbol{\eta} = (\eta_\varnothing, \boldsymbol{\eta}_{A_1}', \boldsymbol{\eta}_{A_2}', \boldsymbol{\eta}_{A_1A_2}')'$, where $\eta_\varnothing = \log(\sum \boldsymbol{\pi}) = 0$ is the null contrast and the remaining vectors have elements specified by (1), (2) and (3), respectively. We will refer to (6) as the bivariate ordered logistic model (BOLM). Lapp et al (1998) show how to fit the Dale and Goodman models starting from the framework of a BOLM. Some computational problems may arise when fitting a multivariate logistic model, depending on the number of responses

and categories. For example, when inverting equation $\boldsymbol{\eta} = \mathbf{C}'\log(\mathbf{L}\boldsymbol{\pi})$ to obtain $\boldsymbol{\pi}$ in terms of $\boldsymbol{\eta}$, it may happen that for certain fixed values of $\boldsymbol{\eta}$ no positive solution $\boldsymbol{\pi}$ exists. Although $\boldsymbol{\pi} > 0$ ensures the matrix $\mathbf{C}'\mathbf{D}^{-1}\mathbf{L}$ to be invertible (Glonek and McCullagh, 1995, Theorem 1), where $\mathbf{D} = \text{diag}(\mathbf{L}\boldsymbol{\pi})$, the range of the mapping is not a hyper-rectangle and fixing some components of $\boldsymbol{\eta}$ restricts the range of the remaining components, that is the model is not *variation independent*. Although this problem is particularly magnified for $K > 2$ responses (Bergsma and Rudas, 2002; Qaqish and Ivanova, 2006), computational problems can also arise in the bivariate case, above all when considering certain particular model configurations. For instance, it may happen that not only the intercepts but some covariates have a category-dependent effect. To highlight this effect one may want to fit a bivariate version of the partial proportional odds model proposed by Peterson and Harrell (1990). However, such a model can be computationally very hard to fit, even with a limited and reasonable number of parameters. To deal with this difficulty, we propose to regularize the parameter space by penalizing the log-likelihood of the model. This allows to increase the range of possible models to be fitted. The penalty term we use for this is introduced in Section 2.1.

The fit of a BOLM becomes computationally hard also when the number of response categories increases. In addition, the model may result overparameterized. Lapp et al (1998) fit a Dale's model by imposing constraints on the row and column interactions of the association intercepts in order to reduce the number of parameters. However, this type of data appears to be too "rich" to be modeled with fully parametric models and nonparametric or semiparametric models, followed by graphical presentation, could result more useful (Eilers and Marx, 1996). In order to smooth the marginal and association effects across the response categories, we suggest to use a penalty term, introduced is Section 2.2, for nonparametric modeling, mainly employed in the P-spline context (Eilers et al, 2006), but suitably rewritten to be used in the framework of a BOLM. In part, this approach can be considered the bivariate extension of the models proposed by Tutz (2003).

The ordinal nature of the responses imposes inequality constraints on marginal distributions which have to be taken into account in model estimation. In Section 2.3, we present a penalty term, which is able to mimic such inequality constraints.

In order to better understand the potential of the penalization approach, some further notation is needed, according to that used in Tutz and Scholz (2003) for the univariate cumulative logistic regression model. Let $\mathcal{Q}$ be the set of indices of all the covariates, excluding the intercepts, and $\mathcal{P} \subset \mathcal{Q}$ be a subset of p covariates. Let $\mathcal{S}$ be the set of indices of the variables whose effects we assume do not depend on categories and such that $\mathcal{S} \subseteq \mathcal{P}$, and let $\bar{\mathcal{S}} = \mathcal{P} \backslash \mathcal{S}$. In particular, we define $\mathcal{S}$ and $\bar{\mathcal{S}}$ as $\mathcal{S} = \cup_{k=1}^{3} \mathcal{S}_k$, and $\bar{\mathcal{S}} = \cup_{k=1}^{3} \bar{\mathcal{S}}_k$, where $\mathcal{S}_k$ and $\bar{\mathcal{S}}_k$ are the subsets of $\mathcal{S}$ and $\bar{\mathcal{S}}$, respectively, associated to the kth equation. To complete the notation, let $\mathcal{S}^0 = \{0\} \cup \mathcal{S}$ and $\bar{\mathcal{S}}^0 = \{0\} \cup \bar{\mathcal{S}}$. Consider the following model where only a part of the covariates is supposed to be category-independent:

$$\begin{cases} \log[\phi_{1r}(\boldsymbol{x}_i)] = \beta_{10r} + \boldsymbol{\beta}'_{1\mathcal{S}_1}\boldsymbol{x}_{i\mathcal{S}_1} + \boldsymbol{\beta}'_{1\bar{\mathcal{S}}_1 r}\boldsymbol{x}_{i\bar{\mathcal{S}}_1}, \\ \log[\phi_{2c}(\boldsymbol{x}_i)] = \beta_{20c} + \boldsymbol{\beta}'_{2\mathcal{S}_2}\boldsymbol{x}_{i\mathcal{S}_2} + \boldsymbol{\beta}'_{2\bar{\mathcal{S}}_2 c}\boldsymbol{x}_{i\bar{\mathcal{S}}_2}, \\ \log[\psi_{rc}(\boldsymbol{x}_i)] = \beta_{30rc} + \boldsymbol{\beta}'_{3\mathcal{S}_3}\boldsymbol{x}_{i\mathcal{S}_3} + \boldsymbol{\beta}'_{3\bar{\mathcal{S}}_3 rc}\boldsymbol{x}_{i\bar{\mathcal{S}}_3}, \end{cases} \tag{7}$$

$(r = 1,\ldots,D_1-1,\ c = 1,\ldots,D_2-1)$. We refer to model (7) as the *Non-Uniform association and Partially Proportional Odds Model* (NUPPOM). Although in the univariate case the phrase "proportional odds" is usually referred to a model with covariate effects which do not depend on the categories, here we will refer to a *Uniform association and Proportional Odds Model* (UPOM) as a model defined from (7) assuming $\beta_{30rc} = \beta_{30}$ and $\bar{\mathcal{S}} = \varnothing$. On the other hand, a *Non-Uniform association and Non-Proportional Odds Model* (NUNPOM) will be defined from (7) assuming $\mathcal{S} = \varnothing$ and with category-dependent association intercepts. Note that the intercepts for the marginal equations (that is the global-logit intercepts) are never supposed to be independent of the categories, whatever the model. According to these definitions the bivariate Dale model (5) is a NUPOM, and it becomes a UPOM when $\rho_{1r} = 0$, $\rho_{2c} = 0$ and $\sigma_{rc} = 0$, $r = 1,\ldots,D_1 - 1$, $c = 1,\ldots,D_2 - 1$. Further, to specify that a NUPPOM is fitted we will also write $NUPPOM(\mathcal{S})$ and to indicate that a UPOM is fitted we will also write $UPOM(\bar{\mathcal{S}}^0)$.

Under multinomial sampling with frequencies $\boldsymbol{y}_i \sim M(n_i, \boldsymbol{\pi}_i)$, consider the model $\mathbf{C}'\log(\mathbf{L}\boldsymbol{\pi}_i) = \mathbf{X}_i\boldsymbol{\beta}$, with the matrices $\mathbf{C}$ and $\mathbf{L}$ such that the marginal parameters are *global logits*, the association parameters are *log global odds ratios*, and the constraint $\sum_{j=1}^{D_1}\sum_{k=1}^{D_2}\pi_{ijk} = 1$ is included. Then, the kernel of the log-likelihood is

$$l(\boldsymbol{\beta}) = \sum_{i=1}^{m} l(\boldsymbol{\beta};\boldsymbol{y}_i) = \sum_{i=1}^{m} \boldsymbol{y}'_i \log(\boldsymbol{\pi}_i), \tag{8}$$

where m, the observed number of response configurations, is such that $\sum_{i=1}^{m} n_i = n$, with n indicating the sample size. The *penalized log-likelihood* has the form

$$l_P(\boldsymbol{\beta}) = l(\boldsymbol{\beta}) - \frac{1}{2}\tau(\boldsymbol{\beta}), \tag{9}$$

where $\tau(\boldsymbol{\beta}) = \boldsymbol{\beta}'\mathbf{P}\boldsymbol{\beta}$, and $\mathbf{P}$ represents the penalization and includes the smoothing parameter. Penalized ML estimation formulas are given and discussed in Appendix A. The specification of $\tau(\boldsymbol{\beta})$ is discussed in the following sections, whereas the form of $\mathbf{P}$ is given in Appendix B.

2.1 A penalty term for parameter space regularization

When the cross-tabulation of the responses contains one or more zeros, parameter estimation by Fisher-scoring may be challenging at each iteration. In these cases, one may try to reduce l_{step}, the step length (see Appendix A). However, estimates of the association structure may result too irregular, with very high (or very low) estimated odds ratios in correspondence of the zero cells. In order to stabilize the ML estimates of the BOLM using Fisher scoring, a reduction of the parameter space may be helpful. We propose to penalize both the marginal and the association parameters. In addition, since the model is not variation independent, applying a penalty term on association parameters might be useful to limit the range of the possible values that the marginal parameters can assume, so avoiding a failure of the Fisher scoring. The general expression is:

$$\tau(\boldsymbol{\beta}) = \sum_{j\in\bar{\mathcal{S}}_1^0}\lambda_{1j}\sum_{r=2}^{D_1-1}\zeta(\beta_{1jr}) + \sum_{j\in\bar{\mathcal{S}}_2^0}\lambda_{2j}\sum_{c=2}^{D_2-1}\zeta(\beta_{2jc}) + \sum_{j\in\bar{\mathcal{S}}_3^0}\lambda_{3j}\sum_{r=2}^{D_1-1}\sum_{c=2}^{D_2-1}\zeta(\beta_{3jrc}), \tag{10}$$

where λ_{kj} is the smoothing parameter for the jth variable of the kth equation of system (7), $k = 1,2,3$. A first specification of $\zeta(.)$, that we call ARC1, is $\zeta(\alpha_t) = (\Delta\alpha_t)^2$, where Δ is the order 1 difference operator, that is $\Delta\alpha_t = \alpha_t - \alpha_{t-1}, t \geq 2$. With respect to (10), the operator acts over the indices r and/or c. It involves penalization of adjacent row and column parameters and is aimed at (*i*) overcoming estimation problems by reducing parameter space and (*ii*) reproducing a *UPOM* for high smoothing values. For the aim (i), the choice of λ is based on the minimum value for which Fisher-scoring does not fail. The simulation study in Section 4 will clarify this choice. The aim (*ii*) is achieved as $\lambda_{kj} \to \infty, k = 1,2,3, \forall j \in \bar{\mathcal{S}}_k^0$, for which all the parameters indexed by j will tend to be equal among the categories. Although (10) allows to penalize marginal intercepts, it is preferable to avoid a strong penalization on such parameters, in order not to violate (12). In this work emphasis is on ARC1, but several other specifications of $\zeta(.)$ are possible. For example, another specification aimed at reducing the parameter space is $\zeta(\alpha) = \alpha^2$, which corresponds to a ridge-type penalty for the bivariate logistic regression model. As $\lambda_{kj} \to \infty, k = 1,2,3, \forall j \in \bar{\mathcal{S}}_k^0$, all the parameters indexed by j will tend to zero. A similar penalty, not involving the third term, is used by Desantis et al (2008) in a penalized latent class model for ordinal data.

2.2 A penalty term for nonparametric modeling

Beside being useful for reducing the parameter space and for reproducing a *UPOM*, the following generalization of the penalty term ARC1, hereafter denoted by ARC2, can be used to specify row or column effects and to fit nonparametric models where the effects are determined by a polynomial:

$$\begin{aligned}\tau(\boldsymbol{\beta}) = &\sum_{j\in\bar{\mathcal{S}}_1^0}\lambda_{1j}\sum_{r=s_1+1}^{D_1-1}(\Delta^{s_{1j}}\beta_{1jr})^2 + \sum_{j\in\bar{\mathcal{S}}_2^0}\lambda_{2j}\sum_{c=s_2+1}^{D_2-1}(\Delta^{s_{2j}}\beta_{2jc})^2 \\ &+\sum_{j\in\bar{\mathcal{S}}_3^0}\left[\lambda_{3j}\sum_{r=s_3+1}^{D_1-1}\sum_{c=1}^{D_2-1}(\Delta^{s_{3j}}\beta_{3jrc})^2\right. \\ &\left. + \lambda_{4j}\sum_{r=1}^{D_1-1}\sum_{c=s_4+1}^{D_2-1}(\Delta^{s_{4j}}\beta_{3jrc})^2\right],\end{aligned} \tag{11}$$

where $\Delta^a = \Delta(\Delta^{a-1})$. Consider the following penalty settings:

- as $\lambda_{hj} = 0, h = 1,...,4, \forall j \in \bar{\mathcal{S}}^0$, an unrestricted model will be fitted;
- as $\lambda_{hj} \to \infty, h = 1,...,4, \forall j \in \bar{\mathcal{S}}_1 \cup \bar{\mathcal{S}}_2 \cup \bar{\mathcal{S}}_3^0$ and $s_{hj} = 1$, the fitted parameters will tend to be equal, and the model will tend to a *UPOM*;
- as $\lambda_{3j} \to \infty$, $\lambda_{4j} = 0, \forall j \in \bar{\mathcal{S}}_3^0$ and $s_{3j} = 1$, a model with column effects will be fitted;
- as $\lambda_{3j} = 0$, $\lambda_{4j} \to \infty, \forall j \in \bar{\mathcal{S}}_3^0$ and $s_{4j} = 1$, a model with row effects will be fitted;
- as $\lambda_{hj} \to \infty, h = 1,2, \forall j \in \bar{\mathcal{S}}^0$ and $s_{hj} > 1$, the fitted parameters will follow a polynomial curve of degree $s_{hj} - 1$.
- as $\lambda_{hj} \to \infty, h = 3,4, \forall j \in \bar{\mathcal{S}}^0$ and $s_{hj} > 1$, the fitted parameters will follow a polynomial surface of degree $s_{3j} + s_{4j} - 2$.

Notice the difference between the penalty terms included in (11) and those included in the penalized log-likelihood (14) in Tutz (2003), suggested for a single ordered response. In that paper, the author proposed to penalize the differences of adjacent categories, for a vertical smoothing, jointly to the use of penalized B-splines for a horizontal smoothing, resulting in a form similar to (11). Also notice the differences with the bivariate horizontal smoothing approach by Bustami et al (2001) which presented the additive bivariate Dale model, for continuous, category-independent covariates, as a natural extension of the generalized additive model (Hastie and Tibshirani, 1990).

Penalty (11) may be useful to assume certain dependence structures on the categories, for both marginal and association parameters. For example, if one wants to assume a linear trend for the row marginal effects, one may assume $\eta_{1ir} = \beta_{10r} + \sum_{j=1}^{p} x_{ij}\beta_{1jr}$, where $\beta_{1jr} = \alpha_{0j} + \alpha_{1j}\delta_{jr}$, with

α_{0j} and α_{1j} unknown parameters, and with scores δ_{jr}. In spite of its simplicity, such an approach assumes arbitrary scores. An alternative way is just to use a penalization approach with penalty term ARC2 which uses scores "chosen by the data" (Tutz and Scholz, 2003). Indeed, the smoothing parameters and the polynomial degrees can be chosen on the basis of some criterion, such as the values that minimize the AIC. As a special case, suppose to want to fit a model which assumes a linear trend of the marginal parameters and an association structure composed by the interaction of two first degree polynomials[1]. By assuming, for simplicity, that the same variable x_j is present in all the equations of system (7), choosing $s_{hj} = 2$, $h = 1,\ldots,4$ and high smoothing values, for instance 10^8, the predictor becomes

- $x_{ij}\beta_{jr} = x_{ij}\gamma_{01j} + x_{ij}(\gamma_{11j}\cdot r)$
- $x_{ij}\beta_{jc} = x_j\gamma_{02j} + x_{ij}(\gamma_{12j}\cdot c)$
- $x_{ij}\beta_{jrc} = x_{ij}\gamma_{03j} + x_j(\gamma_{13j}\cdot r) + x_{ij}(\gamma_{23j}\cdot c) + x_{ij}(\gamma_{33j}\cdot r\cdot c)$,

with scores $\delta_{jr} = r$, $\delta_{jc} = c$, and $\boldsymbol{\delta}_{jrc} = (r,c)'$, that is preassigned equally-spaced scores.

2.3 Mimicking inequality constraints

The ordinal nature of the responses introduces some explicit ordering constraints on marginal distribution which have to be taken into account to avoid ill-conditioning of the predictor space. In particular, for the ith individual, such constraints are on the marginal predictors, that is

$$\beta_{k01}+\boldsymbol{\beta}'_{k1}\boldsymbol{x}_i < \beta_{k02}+\boldsymbol{\beta}'_{k2}\boldsymbol{x}_i < \ldots < \beta_{k0,D_k-1}+\boldsymbol{\beta}'_{k,D_k-1}\boldsymbol{x}_i, \quad (12)$$

$k = 1,2$. Although Lagrangians can be used to take into account such constraints, in the spirit of this paper, a penalized-oriented solution could be the following:

$$\tau(\boldsymbol{\beta}) = \sum_{i=1}^{n}\left[\sum_{k=1}^{2}\lambda_k\sum_{r=2}^{D_k-1} I(\Delta\eta_{kir})(\Delta\eta_{kir})^2\right], \quad (13)$$

where $\eta_{kir} = \beta_{k0r} + \boldsymbol{\beta}'_{kr}\boldsymbol{x}_i$, $\Delta\eta_{kir} = \eta_{kir} - \eta_{ki,r-1}$, and $I(z) = 1$ if $z \geq 0$, otherwise $I(z) = 0$. As $\lambda_k \to \infty$, the penalty term (13) acts in such a way to satisfy (12).The univariate version of (13) is used, for example, by Muggeo and Ferrara (2008) in a penalized splines context applied to univariate generalized linear models. It can also be used jointly to (10) or (11). Notice that, although seemingly superfluous, the inclusion of $(\Delta\eta_{kir})^2$ in (13) derives from the necessity of writing $\tau(\boldsymbol{\beta})$ as a quadratic form in order to exploit the penalized ML formulae in Appendix A.

[1] The degree of a two-variable polynomial is defined as the highest degree of its terms, and the degree of a term is the sum of the exponents of the variables that appear in it. Since (11) allows to fit only polynomial models with interactions, to distinguish each of the possible models having the same degree, it is more practical for us to indicate a model by specifying both the degrees of the one-variable polynomials, omitting to specify the (implicit) presence of interaction terms.

3 Hypothesis testing

When estimates are penalized, the asymptotic distribution of the penalized likelihood ratio (LR_p) statistic is known only for some hypothesis systems. Unfortunately, to check the hypothesis $(P)POM$ to which we are mainly interested, i.e. that of category-independent effects, as far as we know, neither exact nor asymptotic results are known for LR_P. Thus, by a simulation study, we analyze the conditions under which it is possible to approximate, under the hypothesis $(P)POM$, the LR_p asymptotic distribution by a χ^2 distribution. As an introduction, in the following Section 3.1 we first recall a result already present in the literature, which is useful for a simple hypothesis system, and then we show the rationale of using the χ^2 distribution and the results of the simulation study in Section 3.2. To simplify notation we will suppose, without loss of generality, that the same index j refers to the same variable for both marginal and association equations.

3.1 The LR_P statistic for the hypothesis of null effects

Let us consider the specific partition of parameters $\boldsymbol{\beta}_{\mathcal{P}^0} = (\boldsymbol{\gamma},\boldsymbol{\delta})'$, such that the null hypothesis:

$$H_0 : \boldsymbol{\delta} = \mathbf{0}, \quad (14)$$

postulates that only a subset of parameters is constrained. Furthermore, consider the penalized log-likelihood of the more general model, $l_P(\hat{\boldsymbol{\gamma}},\hat{\boldsymbol{\delta}})$, that of the reduced model, $l_P(\tilde{\boldsymbol{\gamma}},\mathbf{0})$ and the penalized log-likelihood ratio statistic:

$$LR_P = -2\{l_P(\tilde{\boldsymbol{\gamma}},\mathbf{0}) - l_P(\hat{\boldsymbol{\gamma}},\hat{\boldsymbol{\delta}})\}. \quad (15)$$

Let $\mathbf{F}$ be the information matrix from the unpenalized partial likelihood, with subscripts denoting the submatrices, such as $\mathbf{F}_{\boldsymbol{\delta\delta}}$ for derivatives with respect to $\boldsymbol{\delta}$. Consider the matrix $\mathbf{F}_{\boldsymbol{\delta\delta}|\boldsymbol{\gamma}} = \mathbf{F}_{\boldsymbol{\delta\delta}} - \mathbf{F}_{\boldsymbol{\delta\gamma}}\mathbf{F}^{-1}_{\boldsymbol{\gamma\gamma}}\mathbf{F}_{\boldsymbol{\gamma\delta}}$. Then, under the null hypothesis, Gray (1994) shows the statistic LR_P to have the same asymptotic distribution as $\sum\alpha_j Z_j^2$, where the Z_j's are independent standard Normal random variables, and the α_j's are the eigenvalues of the matrix $\lim_{n\to\infty}\mathbf{F}_{\boldsymbol{\delta\delta}|\boldsymbol{\gamma}}(\mathbf{F}_{\boldsymbol{\delta\delta}|\boldsymbol{\gamma}}+\mathbf{P})^{-1}$, where $\mathbf{P}$ is the matrix representing the penalty term.

3.2 The LR_P statistic for the $(P)POM$ hypothesis

Consider a full model of the *NUNPOM* type, i.e. for which all variables j, $j \in \mathcal{P}^0 \equiv \{\bar{\mathcal{S}}^0, \mathcal{S} = \varnothing\}$, have category-dependent effects, and a reduced model for which the effects of some variables j, $j \in \mathcal{S} \neq \varnothing$, are category independent. The penalized log-likelihood ratio test to check the hypothesis for comparing these two models, i.e. for testing the null hypothesis:

$$H_0 : \boldsymbol{\beta}_j = \beta_j\mathbf{1},\ j \in \mathcal{S}, \quad (16)$$

compares the maximum penalized log-likelihood $l_P(\hat{\boldsymbol{\beta}}_{\mathcal{P}^0})$, and the maximum penalized log-likelihood $l_P(\tilde{\boldsymbol{\beta}}_{\mathcal{S}}, \tilde{\boldsymbol{\beta}}_{\bar{\mathcal{S}}^0})$:

$$\begin{aligned} LR_P &= -2\{l_P(\tilde{\boldsymbol{\beta}}_{\mathcal{S}}, \tilde{\boldsymbol{\beta}}_{\bar{\mathcal{S}}^0}) - l_P(\hat{\boldsymbol{\beta}}_{\mathcal{P}^0})\} \\ &= 2\sum_{i=1}^{m}\sum_{r=1}^{D_1}\sum_{c=1}^{D_2} \boldsymbol{y}'_{irc} \log\left(\frac{\hat{\boldsymbol{\pi}}_{irc}}{\tilde{\boldsymbol{\pi}}_{irc}}\right) + \tau(\tilde{\boldsymbol{\beta}}) - \tau(\hat{\boldsymbol{\beta}}), \end{aligned} \tag{17}$$

where $\hat{\boldsymbol{\pi}}'_i = (\hat{\pi}_{i11}, \ldots, \hat{\pi}_{iD_1D_2})'$ is the estimated (by penalization) probability vector for the model under H_1 and $\tilde{\boldsymbol{\pi}}_i = (\tilde{\pi}_{i11}, \ldots, \tilde{\pi}_{iD_1D_2})'$ is the corresponding estimated (by penalization) probability vector for the reduced model. Supposing to use the penalty term ARC1 and following Tutz and Scholz (2003), let $\lambda_{kj\mathcal{R}}\,(\lambda_{kj\mathcal{F}})$ denote the smoothing parameters for the reduced model (full model). Then, we have

$$\begin{aligned} &\tau(\tilde{\boldsymbol{\beta}}) - \tau(\hat{\boldsymbol{\beta}}) = \\ &= \sum_{j\in\bar{\mathcal{S}}^0}\Bigg[\lambda_{1j\mathcal{R}}\sum_{r=2}^{D_1-1}(\Delta\tilde{\beta}_{1jr})^2 + \lambda_{2j\mathcal{R}}\sum_{c=2}^{D_2-1}(\Delta\tilde{\beta}_{2jc})^2 \\ &+ \lambda_{3j\mathcal{R}}\sum_{r=2}^{D_1-1}\sum_{c=2}^{D_2-1}(\Delta\tilde{\beta}_{3jrc})^2\Bigg] - \sum_{j\in\mathcal{P}^0}\Bigg[\lambda_{1j\mathcal{F}}\sum_{r=2}^{D_1-1}(\Delta\hat{\beta}_{1jr})^2 \\ &+ \lambda_{2j\mathcal{F}}\sum_{c=2}^{D_2-1}(\Delta\hat{\beta}_{2jc})^2 + \lambda_{3j\mathcal{F}}\sum_{r=2}^{D_1-1}\sum_{c=2}^{D_2-1}(\Delta\hat{\beta}_{3jrc})^2\Bigg] \\ &= \sum_{j\in\bar{\mathcal{S}}^0}\Bigg\{\sum_{r=2}^{D_1-1}\left[\lambda_{1j\mathcal{R}}(\Delta\tilde{\beta}_{1jr})^2 - \lambda_{1j\mathcal{F}}(\Delta\hat{\beta}_{1jr})^2\right] \\ &+ \sum_{c=2}^{D_2-1}\left[\lambda_{2j\mathcal{R}}(\Delta\tilde{\beta}_{2jc})^2 - \lambda_{2j\mathcal{F}}(\Delta\hat{\beta}_{2jc})^2\right] \\ &+ \sum_{r=2}^{D_1-1}\sum_{c=2}^{D_2-1}\left[\lambda_{3j\mathcal{R}}(\Delta\tilde{\beta}_{3jrc})^2 - \lambda_{3j\mathcal{F}}(\Delta\hat{\beta}_{3jrc})^2\right]\Bigg\} \\ &- \sum_{j\in\mathcal{S}}\Bigg[\lambda_{1j\mathcal{F}}\sum_{r=2}^{D_1-1}(\Delta\hat{\beta}_{kjr})^2 + \lambda_{2j\mathcal{F}}\sum_{c=2}^{D_2-1}(\Delta\hat{\beta}_{kjc})^2 \\ &+ \lambda_{3j\mathcal{F}}\sum_{r=2}^{D_1-1}\sum_{c=2}^{D_2-1}(\Delta\hat{\beta}_{kjrc})^2\Bigg]. \end{aligned}$$

If estimates are not penalized, that is if for $k = 1,2,3$, $\lambda_{k1\mathcal{R}} = \lambda_{k2\mathcal{R}} = \cdots = \lambda_{k|\bar{\mathcal{S}}^0|\mathcal{R}} = \lambda_{k0\mathcal{F}} = \lambda_{k1\mathcal{F}} = \cdots = \lambda_{k|\mathcal{P}^0|\mathcal{F}} = 0$, one obtains $v(\tilde{\boldsymbol{\beta}}) - v(\hat{\boldsymbol{\beta}}) = 0$, and the LR_P statistic has the usual asymptotic χ^2 distribution. If $\lambda_{kj\mathcal{R}} = \lambda_{kj\mathcal{F}}$ is chosen for $j \in \bar{\mathcal{S}}^0$, $k = 1,2,3$, then the first term is very small since $\tilde{\beta}_{kjr} \approx \hat{\beta}_{kjr}$, $\tilde{\beta}_{kjc} \approx \hat{\beta}_{kjc}$ and $\tilde{\beta}_{kjrc} \approx \hat{\beta}_{kjrc}$ for $r = 1,\ldots,D_1-1$, $c = 1,\ldots,D_2-1$. Thus, the fundamental term concerns the variables for which $j \in \mathcal{S}$ and, if estimates are penalized with a low smoothing value, converging to zero at an appropriate rate, the same asymptotic behaviour holds.

We show this approximate result by simulation. We simulated the sampling distribution of LR_P assuming to have responses A_1 and A_2, with levels $D_1 = D_2 = 3$, and a single dichotomous covariate X_1, sampled from $Ber(0.5)$. The chosen regression parameters for the true model are the following:

- $\boldsymbol{\beta}_{10} = (-0.5, 0.5)'$, $\boldsymbol{\beta}_{11} = (-0.3, 0.3)'$,
- $\boldsymbol{\beta}_{20} = (-0.1, 0.6)'$, $\boldsymbol{\beta}_{21} = (-0.2, 0.4)'$,
- $\boldsymbol{\beta}_{30} = (1.5, 2, 2.5, 3)'$, $\boldsymbol{\beta}_{31} = (-.5, -.5, -.5, -.5)'$.

The null hypothesis is that of global association effect for variable X_1, that is $H_{\mathcal{S}_3} : \beta_{311} = \beta_{312} = \beta_{313} = \beta_{314} = -.5$. We generated 1500 pseudo-samples of size $n = 400$ from a multinomial distribution with probability matrix $\boldsymbol{\Pi}(\mathbf{X}\boldsymbol{\beta})$ and fitted the reduced and unreduced models by penalizing the association intercepts $\boldsymbol{\beta}_{30}$ according to smoothing values $\lambda_{30} \in \{0, 1, 10, 50\}$. The sampling distribution of LR_P, for varying smoothing parameters, is displayed in Figure 1 together with the superimposed theoretical χ^2_3 distribution, and two vertical lines highlighting the nominal 5% level and the observed level. Observe that the theoretical χ^2_3 approximation to the empirical LR_P distribution is very good for $\lambda_{30} \leq 1$, whereas it is totally wrong for $\lambda_{30} = 50$.

4 Evaluating the performance of penalized estimates

In order to evaluate the potential of smoothed estimates, a small simulation study is carried out. The sample size is taken to be $n = 400$, whereas the number of samples is taken to be $N = 100$. A *NUNPOM* is assumed to hold, with responses A_1 and A_2 with $D_1 = D_2 = 3$ levels, and with

- $\boldsymbol{\beta}_{10} = (-0.6, 0.6)'$, $\boldsymbol{\beta}_{11} = (0.3, -0.3)'$,
- $\boldsymbol{\beta}_{20} = (-0.6, 0.6)'$, $\boldsymbol{\beta}_{21} = (-0.6, 0.6)'$,
- $\boldsymbol{\beta}_{30} = (2.6, 2.4, 2.0, 1.7)'$, $\boldsymbol{\beta}_{31} = (-0.4, 0.2, -0.5, 0.5)'$.

The covariate values x_i, $i = 1,\ldots,n$ were drawn from a uniform distribution on $(-1, 1)$. Thus, given the model formula (6) and the inversion method (4) we found the $n \times (D_1D_2)$ probability matrix $\boldsymbol{\Pi}$, in which each row $\boldsymbol{\pi}'_i$ represents the probability vector for the ith observation. Then, each of n bivariate responses was drawn from a multinomial distribution with probability vector $\boldsymbol{\pi}'_i$. The *UPOM* was compared to the *NUNPOM*, for which ARC1 has been used in combination with (13). Penalization parameters for ARC1, that is λ_1, λ_2 and λ_3 were chosen equal to a single value λ, varying in the set $\{0, 1, 10, 100, \ldots\}$. By starting from the lower λ value, the procedure was stopped when Fisher scoring did not fail. This procedure was iterated for each of N samples. Comparisons were made by evaluating the following loss functions:

Mean squared error loss:

$$MSEL = \frac{1}{n}\sum_{i=1}^{n}\sum_{r=1}^{D_1D_2}(\pi_{ir} - \hat{\pi}_{ir})^2, \tag{18}$$

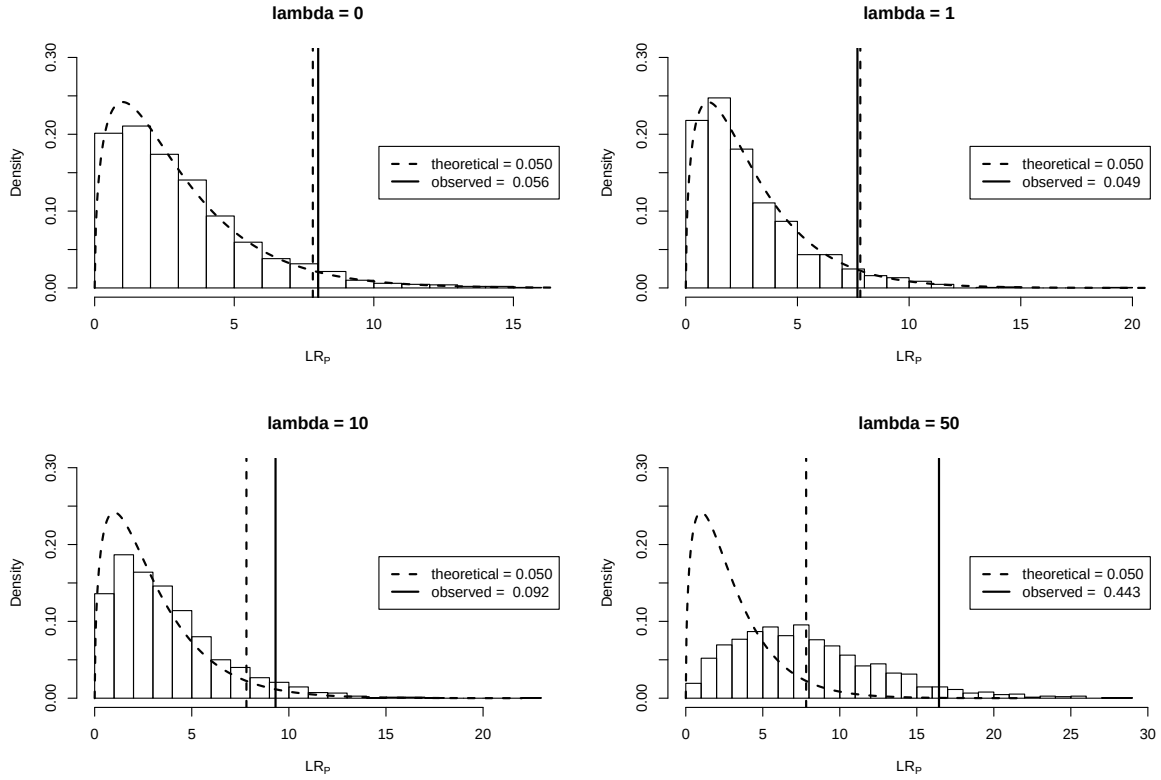

Fig. 1 Simulated distribution of LR_P for varying smoothing parameter and a sample size of 400 simulated observations.

Mean relative squared error loss:

$$MRSEL = \frac{1}{n}\sum_{i=1}^{n}\sum_{r=1}^{D_1D_2}\frac{(\pi_{ir}-\hat{\pi}_{ir})^2}{\pi_{ir}}, \tag{19}$$

Mean entropy or Kullback-Leibler loss:

$$MEL = \frac{1}{n}\sum_{i=1}^{n}\sum_{r=1}^{D_1D_2}\pi_{ir}\log\left(\frac{\pi_{ir}}{\hat{\pi}_{ir}}\right), \tag{20}$$

and the AIC defined in Appendix A. Fisher scoring without penalization failed in 45 out of 100 simulations when the *NUNPOM* was assumed, that is when $\lambda = 0$. Before setting λ to some value greater than zero, some attempts to estimate the model were made by reducing the step length l_{step} of the iterative algorithm (see Appendix A), and in some case the *NUNPOM* was fitted. The *UPOM* was fitted in all the simulations. The results are reported in Table 1.

Observe that for the *NUNPOM* all the mean loss functions and *AIC* are smaller than *UPOM* ones. The *NUNPOM* without penalization (i.e. $\lambda = 0$) has the smallest loss functions values, as it should be, but also the greatest *AIC* value among the *NUNPOMs*. This is due to the trade-off between the best fitting, given by the saturated model, and the greater flexibility, given by the penalization approach, which reduces the degrees of freedom.

Table 1 Comparison *UPOM* vs *NUNPOM* with ARC1, in terms of simulated mean values of the loss functions and *AIC*. *Cumulative number of Fisher scoring successes (*FSS*).

Model	λ	*MSEL*	*MRSEL*	*MEL*	*AIC*	*FSS**
NUNPOM	0	0.0042	0.0408	0.0208	1592.8	55
NUNPOM	1	0.0048	0.0415	0.0225	1579.0	68
NUNPOM	10	0.0046	0.0448	0.0227	1572.2	93
NUNPOM	100	0.0093	0.0927	0.0399	1582.9	100
UPOM	-	0.0129	0.1330	0.0531	1608.8	100

5 Applications to real data sets

5.1 The British males occupational status data set

Consider the data on occupational status (OS) of a sample of British males from Goodman (1979), where fathers and their sons were cross-classified according to the occupational status using seven ordered categories. The data are reported in Table 2.

Several authors have re-analyzed such data. For example, Lapp et al (1998) compare the Goodman RC and Dale models in terms of goodness-of-fit. We further re-analyze the data by fitting the BOLM with ARC2. The aim of the application is to show the advantages of our proposal when compared to the existing alternatives.

Table 2 Cross-classification of British males according to occupational status.

Father's status	Subject's status						
	1	2	3	4	5	6	7
1	50	19	26	8	18	6	2
2	16	40	34	18	31	8	3
3	12	35	65	66	123	23	21
4	11	20	58	110	223	64	32
5	14	36	114	185	714	258	189
6	0	6	19	40	179	143	71
7	0	3	14	32	141	91	106

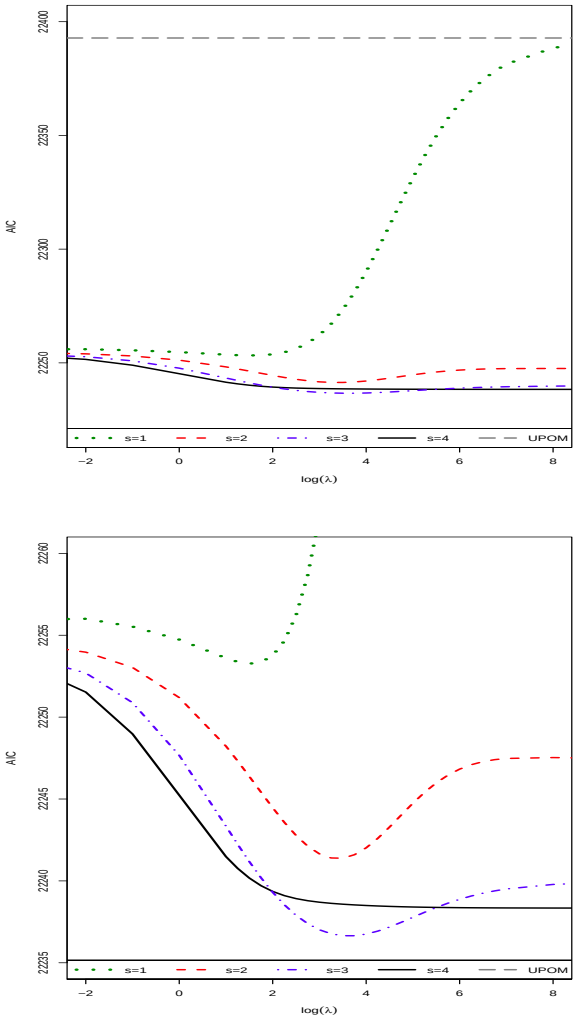

Fig. 2 OS data set, the top graph: AIC for the NUPOM using the penalty term ARC2 for varying smoothing parameter and different orders of penalization. The bottom graph shows a detail of the most critical interval (-2,8).

The saturated model for the joint distribution involves 48 parameters: 6 global logits for each marginal and 36 log-GORs. Since the interest is in modeling the association structure, we concentrate only on the 36 log-GORs. Figure 2 shows the AIC for the NUPOM for varying smoothing parameter and different orders of penalization.

Due to the symmetry of the association structure, the difference operator orders s_3 and s_4 are assumed to be equal and indicated by s. The model with first order penalization, whose *AIC* tends to the UPOM value, is clearly inadequate. The minimum value of *AIC* is 22236.65, corresponding to $\log(\lambda) = 4$ and $s = 3$. This represents a model with a smoothed association structure which tends to a surface defined by row and column interactions of second degree polynomials. Thus, on the grounds of *AIC* only, one could choose this model. However, observe that for high values of λ, the models with $s = 3$ ($AIC = 22239.96$) and $s = 4$ ($AIC = 22238.34$) provide good fits as well, with a slight evidence in favor of the latter model, corresponding to a surface of third degree polynomials. The AIC differences of such models, with respect to the minimum AIC, are respectively 3.31 and 1.69, which are quite small (Burnham and Anderson, 2000, p. 48). When it is possible, as in this case, it is preferable to choose a model providing integer and equally-spaced scores, on the grounds of greater interpretability and for the possibility to use classical test statistics whose asymptotic null distributions are well-known. Therefore, we report in Table 3 the results (in terms of *AIC* and deviance G^2) for the four polynomial models evaluated at the largest value of $\log(\lambda) = 15$, along with the independence and saturated models.

Table 3 Model selection based on AIC and deviance G^2 for the OS data set based on a BOLM with penalty term ARC2. The asterisks indicate a non significant difference with the saturated model.

Model	Description	df	AIC	G^2
1	Independence	36	23081.12	897.52
2	Uniform association	35	22392.83	207.22
3	First degree polynomials	32	22247.46	55.85
4	Second degree polynomials	27	22239.96	*38.36
5	Third degree polynomials	20	22238.34	**22.74
6	Saturated	0	22255.60	0.00

*p=0.07, **p=0.3.

Model 1 has been fitted by using the ridge-type penalty term, such that the estimated global log-odds ratios tend to zero for high values of smoothing parameter. Model 4 provides the most parsimonious but yet acceptable fit, with only 9 estimated parameters and p-value = 0.07. Model 5 estimates only 16 parameters, providing a comparable fit ($G^2 = 22.74$), with a not significant difference with the saturated model (p-value = 0.3). This means the ordinal association structure of occupational status can be well fitted by two polynomials of second or third degree, that is a polynomial surface of fourth or sixth degree, respectively. Notice that Models 4 and 5 are more parsimonious than the best

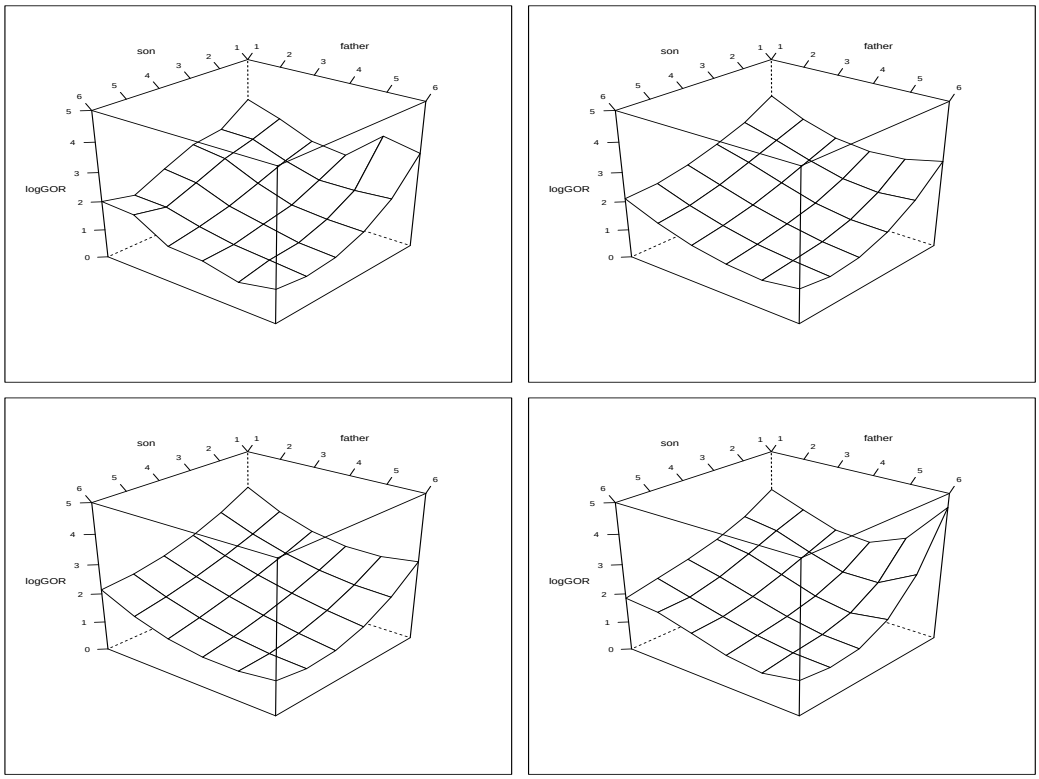

Fig. 3 Association structures for the OS data set. In lexicographical order: the observed (top-left) and fitted log-GORs by the interaction of polynomials of second degree with non-integer scores (top-right), second (bottom-left) and third (bottom-right)degree with integer scores.

model found in Lapp *et al.*'s analysis, i.e. the Dale model, including row effects, column effects, and interactions, while maintaining a comparable fit in terms of G^2. The observed structure of global log-odds ratios and the selected polynomial models are graphically showed in Figure 3.

As we can see, the ordinal association structure is always positive, but it decreases as both the social statuses increase. Observe that the second degree polynomials using integer and non-integer scores show very slight differences. Finally, notice that also Lapp *et al.* hypothesized the possibility to fit a "symmetric second degree polynomial" model.

5.2 The liver disease patients data set

The data set consists of 256 records related to the liver disease progression of patients. The two outcomes, both measured on the same day, are the liver biopsy (named *STAGE*), considered the natural gold standard, and a categorized version of transient elastography (*STIFF*) according to cutoffs suggested by Castera et al (2005), to measure liver stiffness. Both the responses have three ordered categories, *STIFF* with levels 1,2 and 3, which correspond to stiffness classes $[0, 7.1)$, $[7.1, 12.5)$ and $[12.5, \infty)$, respectively; as for *STAGE*,the initial five categories F0-F4 have been collapsed as follows: 1, corresponding to the biopsy stage $< F2$, 2 $(= \{F2, F3\})$ and 3 $(= F4)$. Aim of the study is to evaluate the concordance between the outcomes in order to find profiles of "discordant" patients. This is done using the bivariate logistic model, by employing the log global odds as marginal parameters and the log global odds ratio as association measure. For these data, a first analysis with dichotomized responses was made by Calvaruso et al (2010) but estimation problems inhibited an analysis with 3 levels for each response. Table 4 shows the cross-classification of the responses, ignoring covariates.

Table 4 Marginal Cross-classification of the responses and empirical global log-odds ratios for the liver disease patient data.

	STIFF		
STAGE	1	2	3
1	71	20	0
	(1.72)	$(+\infty)$	
2	56	20	8
	(3.18)	(3.31)	
3	8	27	46

From a first look at Table 4, it is possible to notice an overall positive association between the responses, even if there are many discordant patients, mainly the fifty-six in $STAGE = 2, STIFF = 1$. Among the covariates, the patient's sex (*SEX*), age (*AGE*), alanine aminotransferase (*ALT*) measured in U/L and platelet (*PLT*) levels measured in 10^3 mmc, are considered. For modelling purposes, the covariates were centered with respect to their means and, after a backward selection, the following sets of variables have been considered:

$$\mathcal{P}_1^0 = \{\text{Intercept}_{\overline{123}}, SEX_{123}, AGE_{123}, ALT_{123}, PLT_{\overline{123}}\},$$
$$\mathcal{P}_2^0 = \{\text{Intercept}_{\overline{123}}, SEX_{123}, AGE_{123}, ALT_{123}, PLT_{123}\},$$
$$\mathcal{P}_3^0 = \{\text{Intercept}_{\overline{123}}, SEX_{123}, AGE_{123}, ALT_{123}, PLT_{\overline{123}}\},$$
$$\mathcal{P}_4^0 = \{\text{Intercept}_{\overline{123}}, SEX_{123}, AGE_{123}, ALT_{123}, PLT_{123}\},$$
$$\mathcal{P}_5^0 = \{\text{Intercept}_{\overline{123}}, AGE_{123}, ALT_{123}, PLT_{\overline{123}}\},$$
$$\mathcal{P}_6^0 = \{\text{Intercept}_{\overline{123}}, AGE_{12}, ALT_{123}, PLT_{\overline{12}}\},$$

where $\text{Intercept}_{\overline{123}}$ indicates that the marginal and association intercepts are category dependent, whereas $\text{Intercept}_{\overline{123}}$ indicates that the intercepts for the association are category independent. Here, to indicate that variable *AGE* is included both in marginal and association predictors, we use AGE_{123}, whereas AGE_{12} indicates that such variable is included only in the marginal predictors.

Computational problems have arisen when we tried to estimate the models which assumed a non-uniform association. We have overcome such problems by regularizing the parameter space of the association intercepts. In particular, we have employed the ARC1 penalty term, with smoothing parameter $\lambda_3 = 1$ which is the minimum common penalization value for which Fisher scoring did not fail in all the fitted models of $NU(P)POM$ type. By considering the results from the simulation in Section 3, we decided to use a χ^2 distribution to approximate the LR_P asymptotic distribution. This is done in the model selection reported in Table 5.

For each model we have selected, the table reports its description, the number of estimated parameters and the AIC. The next columns refer to comparisons between nested models, specified by the column headed "vs". The last three columns report the results of such a comparison in terms of penalized log-likelihood ratio statistic, along with degrees of freedom and p-values. Before proceeding to variable selection we first checked, for each variable, the hypothesis $UPOM$, versus several alternatives: $UPPOM$, $NUPOM$, $NUPPOM$. The table reports such comparisons for Models 1-4. Model 1 is the most complex model we have considered, a $NUPPOM$ defined on set $\mathcal{P}_1^0$. This model assumes that the effect of variable *PLT* on $STIFF$ depends on the categories of such response variable. Models 2-4 represent hypotheses of uniform association and/or (partially) proportional odds, and these models are compared to Model 1, for which none of these hypotheses holds. Although the LR_P test for model comparison is approximated, some results seem to be clear. For example, the difference between model 1 and 3 (or 4) is highly significant, so we can claim that the hypothesis of $UPPOM$ (or $UPOM$) does not hold. Models 5-7 concern a backward model selection starting from Model 1. The last row reports the comparison between Models 7 and 1, for which the difference between the starting model and the final model is not significant (p-value=0.711). In model 7, variable *ALT* is the only one which has significant (global) effect for the association model. By AIC, the model with the best trade-off between goodness-of-fit and parsimony is still Model 7. Estimates for this final model are reported in Table 6.

Table 5 Model selection based on the *AIC* and the LR_P statistic for the liver disease patients data using the ARC1 penalty term.

Model	Description	# par.	AIC	vs	LR_P	df	p-value
1	$NUPPOM(\mathcal{P}_1^0)$	21	879.5	-	-	-	-
2	$NUPOM(\mathcal{P}_2^0)$	20	882.1	1	4.67	1	0.031
3	$UPPOM(\mathcal{P}_3^0)$	18	895.7	1	16.51	3	< 0.001
4	$UPOM(\mathcal{P}_4^0)$	17	898.2	1	20.99	4	< 0.001
5	$NUPPOM(\mathcal{P}_5^0)$	18	874.7	1	1.35	3	0.719
6	$NUPPOM(\mathcal{P}_6^0)$	17	873.0	5	0.36	1	0.549
7	$NUPPOM(\mathcal{P}_7^0)$	16	872.7	6	1.23	1	0.268
7	-	-	-	1	2.93	5	0.711

For both outcomes, variables *ALT*, *AGE* and *PLT* are significant. In particular the platelet level has a category-dependent effect for $STIFF$ which is higher for the log global odds 1-2 than 3. In particular, a patient at older age, higher *ALT* and lower *PLT* values is more at risk of having a greater liver stiffness than a patient with mean values. In addition, the *ALT* effect for $STIFF$ is about twice as strong as for biopsy stage. The effect of *ALT* in the association is significant, and considering the intercepts values as well, higher *ALT* values imply a global reduction of the association, especially for individuals in class $STIFF < 7.1$ and $STAGE = 1$.

Table 6 Estimates for Model 4.

response	variable	estimate	se	z	p.value
STAGE	Intercept 1	-0.7398	0.1417	-5.2217	0.0000
	Intercept 2	0.8798	0.1452	6.0610	0.0000
	ALT	-0.0047	0.0018	-2.6217	0.0088
	AGE	-0.0407	0.0105	-3.8985	0.0001
	PLT	0.0093	0.0021	4.4209	0.0000
STIFF	Intercept 1	0.1021	0.1377	0.7421	0.4581
	Intercept 2	1.7861	0.2003	8.9152	0.0000
	ALT	-0.0089	0.0019	-4.6721	0.0000
	AGE	-0.0422	0.0110	-3.8274	0.0001
	PLT 1	0.0087	0.0024	3.6768	0.0002
	PLT 2	0.0154	0.0030	5.1944	0.0000
STAGE vs STIFF	Intercept	1.5164	0.3320	4.5680	0.0000
	Intercept	3.3436	0.6207	5.3869	0.0000
	Intercept	2.8285	0.3880	7.2892	0.0000
	Intercept	2.8134	0.4189	6.7161	0.0000
	ALT	-0.0075	0.0036	-2.0668	0.0387

6 Discussion

We have shown how to fit a BOLM by penalized ML estimation with some penalty terms for a "vertical penalization", that is across response levels. Particular emphasis on the terms ARC1 and ARC2, penalizing adjacent row and column effects, has been given. The motivation for such an approach is, on one hand, its flexibility in modeling situations in which ML estimation by traditional Fisher scoring appears somewhat difficult, and on the other hand, the possibility to consider the fit of a *NUPPOM*, which lies between a *UPOM*, which may give a poor fit, and a *NUNPOM*, often less useful and somewhat more complicated to estimate than a *UPOM*. The penalized log-likelihood ratio LR_P statistic has been considered to check the hypothesis that certain effects are category independent. To our knowledge, the asymptotic distribution of LR_P for the considered hypothesis is not known, though we have shown, by simulation, that for relatively small smoothing values the χ^2 may be a good approximation. However, as far as the distributional properties of penalized likelihood ratio test-statistics are concerned, further investigations are necessary. The potential of penalized estimates by penalty term ARC1 has been shown by simulation and by an application to an original data set. In addition, the BOLM has been fitted using the penalty term ARC2 to a literature data set for comparison with the alternative Dale and Goodman RC models, showing parsimony while preserving a satisfactory the goodness-of-fit. In some sense, ARC2 generalizes ARC1, permitting to fit restricted versions of the Dale model, by inserting row or column effects, but also polynomial effects models, with scores chosen by data. All the computations, including the model implementation, have been carried out by an original R code which can be requested from the authors. Furthermore, an R package implementing all the methods proposed in this paper is under preparation.

Appendix A: penalized maximum likelihood estimation

Let $\partial l/\partial \boldsymbol{\pi}_i = \text{diag}(\boldsymbol{\pi}_i)^{-1}\boldsymbol{y}_i$, $\partial \boldsymbol{\pi}_i/\partial \boldsymbol{\eta}_i = (\mathbf{C}'\mathbf{D}_i^{-1}\mathbf{L})^{-1}$ and $\partial \boldsymbol{\eta}_i/\partial \boldsymbol{\beta} = \mathbf{X}_i$. By using the chain rule, the first derivative of the penalized log likelihood with respect to $\boldsymbol{\beta}$ is

$$\frac{\partial l_P}{\partial \boldsymbol{\beta}} = \sum_{i=1}^{m} \frac{\partial l}{\partial \boldsymbol{\pi}_i}\frac{\partial \boldsymbol{\pi}_i}{\partial \boldsymbol{\eta}_i}\frac{\partial \boldsymbol{\eta}_i}{\partial \boldsymbol{\beta}} - \mathbf{P}\boldsymbol{\beta},$$

the *penalized score function* is

$$\boldsymbol{s}_P(\boldsymbol{\beta};\boldsymbol{y}_i) = \sum_{i=1}^{m} [(\mathbf{C}'\mathbf{D}_i^{-1}\mathbf{L})^{-1}\mathbf{X}_i]'\text{diag}(\boldsymbol{\pi}_i)^{-1}\boldsymbol{y}_i - \mathbf{P}\boldsymbol{\beta},$$

and the *penalized Fisher matrix* is

$$\mathbf{F}_P(\boldsymbol{\beta}) = \sum_{i=1}^{m} n_i\mathbf{X}_i'(\mathbf{L}'\mathbf{D}_i'^{-1}\mathbf{C})^{-1}\text{diag}(\boldsymbol{\pi}_i)^{-1}(\mathbf{C}'\mathbf{D}_i^{-1}\mathbf{L})^{-1}\mathbf{X}_i + \mathbf{P}.$$

Using these formulas, the $(k+1)$th iteration of the Fisher scoring is $\hat{\boldsymbol{\beta}}^{(k+1)} = \hat{\boldsymbol{\beta}}^{(k)} + l_{step}\mathbf{F}_P(\hat{\boldsymbol{\beta}}^{(k)})^{-1}\boldsymbol{s}_P(\hat{\boldsymbol{\beta}}^{(k)})$, where l_{step} is a positive scalar representing the step length. Since the iterative procedure may produce incompatible $\boldsymbol{\beta}$ values for $\boldsymbol{\pi}$, a value smaller than 1 for l_{step}, say 0.5 or smaller, may be necessary, even if this inevitably increases the number of iterations. As a reasonable starting value for $\boldsymbol{\beta}$, one could set to zero the regression coefficients corresponding to covariates, together with the global log-odds ratios intercepts, whereas the global logits intercepts have to be chosen by taking into account the inequality constraints (12). The variance covariance matrix of $\hat{\boldsymbol{\beta}}$ is given by $V(\hat{\boldsymbol{\beta}}) = F_P(\hat{\boldsymbol{\beta}})^{-1}$. When a *NUPPOM* is considered, the form of matrix $\mathbf{X}_i$ is $\mathbf{X}_i = \bigoplus_{k=1}^{3}\mathbf{X}_{k,i}$, where:

$$\mathbf{X}_{k,i} = \begin{pmatrix} 1 & 0\ \boldsymbol{x}_{i,\mathcal{S}_k}' & \boldsymbol{x}_{i,\bar{\mathcal{S}}_k}' & \mathbf{0}' \\ \ddots & \vdots & \ddots & \\ 0 & 1\ \boldsymbol{x}_{i,\mathcal{S}_k}' & \mathbf{0}' & \boldsymbol{x}_{i,\bar{\mathcal{S}}_k}' \end{pmatrix}.$$

Thus the full design matrix is simply $\mathbf{X} = (\mathbf{X}_1', \mathbf{X}_2', ..., \mathbf{X}_m')'$. The *weight function* for the ith observation is defined as $\mathbf{W}_i(\boldsymbol{\beta}) = n_i\left(\frac{\partial \boldsymbol{\pi}_i}{\partial \boldsymbol{\eta}_i'}\right)\text{diag}(\boldsymbol{\pi}_i)^{-1}\left(\frac{\partial \boldsymbol{\pi}_i}{\partial \boldsymbol{\eta}_i}\right)$, the weight matrix is $\mathbf{W}(\boldsymbol{\beta}) = (\mathbf{W}_1(\boldsymbol{\beta})', \mathbf{W}_2(\boldsymbol{\beta})', \ldots, \mathbf{W}_m(\boldsymbol{\beta})')'$, the *hat matrix* is $\mathbf{H} = \mathbf{X}(\mathbf{X}'\mathbf{W}(\hat{\boldsymbol{\beta}})\mathbf{X} + \mathbf{P})^{-1}\mathbf{X}'\mathbf{W}(\hat{\boldsymbol{\beta}})$, and the *Akaike Information Criterion* is $AIC = -2(l(\hat{\boldsymbol{\beta}}) - tr(\mathbf{H}))$.

Appendix B: penalty terms in matrix form

When (10) or (11) is used $\mathbf{P} = \mathbf{E}'\boldsymbol{\Lambda}'^{1/2}\boldsymbol{\Lambda}^{1/2}\mathbf{E}$, where $\boldsymbol{\Lambda}$ is the matrix of smoothing values, and $\mathbf{E} = \bigoplus_{k=1}^{3}\mathbf{E}_k$.

In (10) matrices $\boldsymbol{\Lambda}$ and $\mathbf{E}_k$ depend on the penalty (ridge or ARC1). Let $\boldsymbol{d} = (D_1-1, D_2-1, (D_1-1)(D_2-1))'$ a vector indexed by d_k, $k = 1,2,3$, and let $\boldsymbol{\lambda}'_{k_{|\bar{\mathcal{S}}_k^0|}} = (\lambda_{k,0}, \boldsymbol{\lambda}'_{k_{|\bar{\mathcal{S}}_k|}})$ be the smoothing values vector of length $|\bar{\mathcal{S}}_k^0|$, that is the cardinality of the set of variables undergone to penalization for the kth equation in (7). Then

$$\boldsymbol{\Lambda} = \mathrm{diag}(\boldsymbol{\lambda}'_{1,\mathcal{P}_1^0}, \boldsymbol{\lambda}'_{2,\mathcal{P}_2^0}, \boldsymbol{\lambda}'_{3,\mathcal{P}_3^0}),$$

where $\boldsymbol{\lambda}'_{k,\mathcal{P}_k^0} = (\lambda_{k,0}\mathbf{1}'_{(d_k-1)}, \mathbf{0}'_{|\mathcal{S}_k|}, \boldsymbol{\lambda}'_{k_{|\bar{\mathcal{S}}_k|}} \otimes \mathbf{1}'_{(d_k-1)})$.

Ridge-type penalty term

Let $\mathbf{I}_{d_k}$ be the $d_k \times d_k$ identity matrix. Then

$$\mathbf{E}_k = (\mathbf{I}_{d_k}, \mathbf{0}_{d_k \times |\mathcal{S}_k|}, \mathbf{1}'_{|\bar{\mathcal{S}}_k|} \otimes \mathbf{I}_{d_k}).$$

ARC1 penalty term

For the ARC1 penalty, let $\mathbf{T}_k$ be the $d_k \times d_k$ upper triangular matrix of 1's. Its inverse $\mathbf{V}_k = \mathbf{T}_k^{-1}$ has 1's on the main diagonal, -1's on the first superdiagonal and 0's elsewhere. Further, let $\mathbf{V}_{k_{-1}}$ be the matrix $\mathbf{V}_k$ ignoring the last row. Then

$$\mathbf{E}_k = (\mathbf{V}_{k_{-1}}, \mathbf{0}_{(d_k-1)\times|\mathcal{S}_k|}, \mathbf{1}'_{|\bar{\mathcal{S}}_k|} \otimes \mathbf{V}_{k_{-1}}).$$

ARC2 penalty term

Let $\boldsymbol{c} = (D_1-1, D_2-1, D_2-1, D_1-1)'$ a vector indexed by c_h, $h = 1,\ldots,4$, and let $k = 1,2,3$. Then

$$\boldsymbol{\Lambda} = \mathrm{diag}(\boldsymbol{\lambda}'_{1,\mathcal{P}_1^0}, \boldsymbol{\lambda}'_{2,\mathcal{P}_2^0}, \boldsymbol{\lambda}'_{3,\mathcal{P}_3^0}, \boldsymbol{\lambda}'_{4,\mathcal{P}_3^0}),$$

where $\boldsymbol{\lambda}'_{h,\mathcal{P}_k^0} = (\lambda_{h,0}\mathbf{1}'_{(c_h-1)}, \mathbf{0}'_{|\mathcal{S}_k|}, \mathbf{1}'_{(c_h-1)} \otimes \boldsymbol{\lambda}'_{h_{|\bar{\mathcal{S}}_k|}})$.

Define $s_{h,j}$, $j \in \bar{\mathcal{S}}^0$, the order of operator $\Delta^{s_{h,j}}$, for the jth variable, also including the intercepts. Let $\mathbf{T}_h$ be the $c_h \times c_h$ upper triangular matrix of $1's$ and let $\mathbf{V}_h = \mathbf{T}_h^{-1}$. Let $\mathbf{V}^{s_{h,j}} = \prod_{h=1}^{s_{h,j}} -\mathbf{V}_h$, let $\mathbf{V}^{s_{h,j}}_{-s_{h,j}}$ be the matrix $\mathbf{V}^{s_{h,j}}$ ignoring the last $s_{h,j}$ rows and let $\mathbf{U}_{k,\bar{\mathcal{S}}^0} = (\mathbf{U}_{k,0}, \mathbf{U}_{k,\bar{\mathcal{S}}})$, where $\mathbf{U}_{k,\bar{\mathcal{S}}} = (\mathbf{U}_{k,1},\ldots,\mathbf{U}_{k,|\bar{\mathcal{S}}|})$ and $\mathbf{U}_{1,j} = \mathbf{V}^{s_{1,j}}_{-s_{1,j}}$, $\mathbf{U}_{2,j} = \mathbf{V}^{s_{2,j}}_{-s_{2,j}}$ and $\mathbf{U}_{3,j} = (\mathbf{I}_{(c_3)} \otimes \mathbf{V}^{s_{3,j}}_{-s_{3,j}}) \oplus (\mathbf{V}^{s_{4,j}}_{-s_{4,j}} \otimes \mathbf{I}_{(c_4)})$. Then

$$\mathbf{E}_k = (\mathbf{U}_{k,0}, \mathbf{0}_{(d_k-1)\times|\mathcal{S}_k|}, \mathbf{U}_{k,\bar{\mathcal{S}}}).$$

The penalty term for ordering constraints

For (13) let $\mathbf{N} = (\mathbf{1}_n \otimes \mathbf{E})'$, being $\mathbf{E}$ defined as for ARC1, and n the sample size. Let $\boldsymbol{\Lambda} = (\lambda_1\mathbf{I}_n, \lambda_2\mathbf{I}_n, \mathbf{0}_{n\times n})$. Then

$$\mathbf{P} = \mathbf{X}'\mathbf{N}'\boldsymbol{\Lambda}'^{1/2} I(\boldsymbol{\beta}'\mathbf{X}'\mathbf{N}') I(\mathbf{N}\mathbf{X}\boldsymbol{\beta})\boldsymbol{\Lambda}^{1/2}\mathbf{N}\mathbf{X},$$

where $I(\mathbf{N}\mathbf{X}\boldsymbol{\beta} \leq 0)$ is element-wise, that is $I(a_{ij} \leq 0) = 1$ if true, 0 otherwise.

Acknowledgements This paper has been supported by the Italian Ministerial grant PRIN 2008 "Measures, statistical models and indicators for the assessment of the University System", n. 2008WXSLH.

References

Bartolucci F, Forcina A (2002) Extended RC association models allowing for order restrictions and marginal modeling. J Am Stat Assoc 97(460):1192–1199

Bergsma WP, Rudas T (2002) Marginal models for categorical data. Ann Stat 30:140–159

Burnham K, Anderson D (2000) Model selection and inference. A practical information - Theoretic approach. Springer-Verlag, New York

Bustami R, Lesaffre E, Molenberghs G, Loos R, Danckaerts M, Vlietinck R (2001) Modelling bivariate ordinal responses smoothly with examples from ophthalmology and genetics. Stat Med 20:1825–1842

Calvaruso V, Cammá C, Marco VD, Maimone S, Bronte F, Enea M, Dardanoni V, Manousou P, Pleguezuelo M, Xirouchakis E, Attanasio M, Dusheiko G, Burroughs AK, Craxí A (2010) Fibrosis staging in chronic hepatitis C: analysis of discordance between transient elastography and liver biopsy. J Viral Hepat 17(7):469–474

Castera L, Vergniol J, Foucher J, Bail BL, Chanteloup E, Haaser M, Darrieth M, Couzigou P, Lédinghen VD (2005) Prospective comparison of transient elastography, Fibrotest, APRI and Liver Biopsy for the assessment of fibrosis in chronic hepatitis C. Gastroenterology 138:343–350

Colombi R, Forcina A (2001) Marginal regression models for the analysis of positive association of ordinal response variables. Biometrika 88(4):1007–1019

Dale JR (1986) Global Cross-Ratio Models for Bivariate, Discrete, Ordered Responses. Biometrics 42(4):909–917

Dardanoni V, Forcina A (1998) A unified approach to likelihood inference on stochastic orderings in a nonparametric context. J Am Stat Assoc 93:1112–1123

Desantis SD, Houseman EA, Coull BA, Stammer-Rachamimov A, Betensky RA (2008) A penalized latent class model for ordinal data. Biostatistics 9(2):249–262

Eilers PHC, Marx BD (1996) Flexible smoothing with B-splines and penalties. Stat Sci 11:89–121

Eilers PHC, Currie ID, Durbán M (2006) Fast and compact smoothing on large multidimensional grids. Comput Stat Data An 50:61–76

Fahrmeier L, Gieger C, Heumann C (1999) An application of isotonic longitudinal marginal regression to monitoring the healing process. Biometrics 55:951–956

Gieger C (1997) Non- and semiparametric marginal regression models for ordinal response. Tech. rep., University of Munich, Institute of Statistics, URL `http://epub.ub.uni-muenchen.de/`

Glonek GFV, McCullagh P (1995) Multivariate logistic models. J Roy Stat Soc B Met 57:533–546

Goodman LA (1979) Simple models for the analysis of cross-classifications having ordered categories. J Am Stat Assoc 74(367):537–552

Gray RJ (1994) Spline-based test in Survival Analysis. Biometrics 50:640–652

Hastie TJ, Tibshirani RJ (1990) Generalized Additive Models. Chapman & Hall, London

Lapp K, Molenberghs G, Lesaffre E (1998) Models for the association between ordinal variables. Comput Stat Data An 28:387–411

Muggeo VMR, Ferrara G (2008) Fitting generalized linear models with unspecified link function: A p-spline approach. Comput Stat Data An 52:2529–2537

Peterson B, Harrell FE (1990) Partial proportional odds models for ordinal response variables. J Roy Stat Soc C-App 39(2):pp. 205–217

Plackett RL (1965) A class of bivariate distributions. J Am Stat Assoc 60:516–522

Qaqish BF, Ivanova A (2006) Multivariate logistic models. Biometrika 93(4):1011–1017

Tutz G (2003) Generalized semiparametrically structured ordinal models. Biometrics 59(2):263–273

Tutz G, Scholz T (2003) Ordinal regression modelling between proportional odds and non-proportional odds. Tech. rep., University of Munich, Institute of Statistics, URL `http://epub.ub.uni-muenchen.de/`